Collective Search by Mobile Robots using Alpha-Beta Coordination

Steven Y. Goldsmith and Rush Robinett,III
Sandia National Laboratories[1]
Albuquerque, NM 87185-5800, U.S.A.
Phone: 505-845-8926
sygolds@sandia.gov, rdrobin@isrc.sandia.gov

Keywords: collective robotics, collective search, reactive coordination, emergent behavior, collaborative agents

Abstract. One important application of mobile robots is searching a geographical region to locate the origin of a specific sensible phenomenon. Mapping mine fields, extraterrestrial and undersea exploration, the location of chemical and biological weapons, and the location of explosive devices are just a few potential applications. Teams of "robotic bloodhounds" have a simple common goal; to converge on the location of the source phenomenon, confirm its intensity, and to remain aggregated around it until directed to take some other action. In cases where human intervention through teleoperation is not possible, the robot team must be deployed in a territory without supervision, requiring an autonomous decentralized coordination strategy. This paper presents the *alpha-beta* coordination strategy, a family of collective search algorithms that are based on dynamic partitioning of the robotic team into two complementary social roles according to a sensor-based status measure. Robots in the *alpha* role are risk-takers, motivated to improve their status by exploring new regions of the search space. Robots in the *beta* role are motivated to improve but are conservative, and tend to remain aggregated and stationary until the alpha robots have identified better regions of the search space. Roles are determined dynamically by each member of the team based on the status of the individual robot relative to the current state of the collective. Partitioning the robot team into alpha and beta roles results in a balance between exploration and exploitation, and can yield collective energy savings and improved resistance to sensor noise and defectors. Alpha robots waste energy exploring new territory, and are more sensitive to the effects of ambient noise and to defectors reporting inflated status. Hypothetically, beta robots conserve energy by moving in a direct path to regions of confirmed high status. Beta robots also resist the effects of noise and defectors by averaging status data, but must rely on alpha robots to improve their performance. Alpha-beta is a reactive strategy that requires directed communication of instantaneous sensor data among team members, but does not rely on a domain model. Alpha-beta coordination is a new and ongoing research effort. We present the basic concepts behind the alpha-beta strategy and exhibit preliminary simulation data that illustrate the collective search modes in an idealized search domain.

[1] Sandia is a multiprogram laboratory operated by Sandia Corporation, a Lockheed Martin Company, for the United States Department of Energy under Contract DE-AC04-94AL85000.

1 Introduction and Motivation

Many challenging new applications in robotics involve distributed search and sensing by a robotic team. Mapping mine fields, extraterrestrial and undersea exploration, exploring volcanoes, the location of chemical and biological weapons, and the location of explosive devices are just a few. This paper presents initial but ongoing research into the issues of collective and emergent behaviors in teams of mobile robots tasked with locating specific sensory phenomena. Our motivation for this line of inquiry is the engineering and eventual deployment of large numbers of inexpensive, expendable sensory robots in hazardous or hostile environments, with a particular emphasis on sensing concentrations of hazardous chemicals in terrestrial environments. The problem suite of interest involves the most demanding of sensing environments; rough terrain with obstacles, non-stationary and dilute chemical concentrations, deliberate interference by hostile robots, and limited opportunities for human interaction with the robots through teleoperation [Klarer 1998]. Because human intervention is not always possible in these environments, decentralized coordination schemes which feature collective decision-making by individual autonomous robots are the most promising avenues of research. Overcoming the limitations of crude but inexpensive chemical sensors by using distributed signal processing algorithms that utilize shared data from a large number of agents is another important concept to be investigated. An important issue not generally addressed in robotics research is deliberate and subtle interference with the goals of the robotic team by impostor robots.

Geographical search problems that use robotic teams can be divided into three broad classes: source identification, source mapping, and source localization [Goldsmith and Robinett 1998a]. Robots performing source identification must answer the question "Does region R contain phenomenon X"? A simple yes or no is an adequate answer, and the task can in principle be accomplished by a robot team without actually localizing the target phenomenon. Source mapping requires the robot team to perform an exhaustive search of an area and to localize all phenomena within the region. Source localization problems require precise localization of a target source within a given region. In the simplest form of the source localization problem, a single sensible source is present somewhere within the search space. The search space is divided into two regions based on the quality of sensor data available in the region. The *insensate region* is characterized by a low signal-to-noise ratio. Robots roaming in this region are without information to guide their search activities, and effective search requires multi-agent coordination mechanisms that involve explicit collaboration [Cao, Fukunaga, Kahng, and Meng 1993]. Some coordination strategies for organized collaborative search in zero-information environments are discussed in [Spires & Goldsmith 1998]. The *sensate region* contains the source and is characterized by a signal-to-noise ratio significantly greater than unity. Robots operating in the sensate region have usable but noisy sensory information to guide their search.

Designing a mobile robot team to search a sensate region for a specific target phenomenon involves numerous engineering tradeoffs among robot attributes and environmental variables. For example, battery-driven robots have a finite energy store and can only search a finite area before running down. Success at finding a target source with finite energy resources depends on other characteristics of the robot such as sensor accuracy and noise and efficiency of the locomotive subsystem, as well as properties of the collective such as the number of robots in the team, the use of shared information to reduce redundant search, and the team coordination strategy used to ensure a coherent search process.

2 Alpha-beta Coordination

This paper is concerned with solving the source localization problem using a decentralized coordination strategy we call *alpha-beta coordination.* The *alpha-beta* coordination strategy is a family of collective search algorithms that allow teams of communicating agents[2] to implicitly coordinate their search activities through a division of labor based on self-selected roles and social status. In an alpha-beta team, an agent plays one of two complementary roles. Agents in the *alpha* role are motivated to improve their status by exploring new regions of the search space. Agents in the *beta* role are also motivated to improve, but are conservative and tend to remain aggregated and stationary until the alpha agents have clearly identified better regions of the search space. An agent selects its current role dynamically based on its current status value relative to the current status values of the other team members. Status is determined by some function of the agent's sensor readings, and is generally a measurement of source intensity at the agent's current location. An agent's decision cycle comprises three sequential decision rules: (1) selection of the current role based on the evaluation of the current status data; (2) selection of a specific subset of the current data; and (3) computation of the next heading using the selected data. Variations of these decision rules produce different versions of alpha and beta behaviors that lead to different global properties.

Partitioning the robot team into alpha and beta roles produces a balance between exploration and exploitation. Alpha agents waste energy exploring low-status regions of the search space, but communicate valuable state information to team members that prevents costly reexploration of low-status regions. Alpha agents by nature seek to emulate and ultimately surpass the highest-performing team members and are therefore more sensitive to the effects of transient noise and are more susceptible to the influence defectors[3] reporting false status values. Beta agents use energy wisely by resisting transient influences and moving in a direct path to high-status regions of the search space identified by alpha agents. Hypothetically, beta agents resist noise and defectors (we do not provide support for this claim herein) by selective re-sampling and averaging of status data, but must rely on alpha robots to improve their performance. Consequently, beta agents can be mislead by noise and defectors under some circumstances through second order effects if many of the the alpha agents are mislead. Alpha-beta coordination relies on the following assumptions:

[2] We will use the term agent hereafter to signify the generality of the alpha-beta concept and to stress that we have not yet implemented the technique on actual robotic vehicles.

[3] Defectors may inadvertently misrepresent their status because of flaws, or may be impostors that deliberately attempt to mislead the loyal team members. These kind of effects can be characterized as Byzantine failures [Lamport, Shastak, and Pease 1982].

1. Team members have a reliable communications mechanism.

2. The team is positioned in the (noisy) sensate region surrounding a target source.

3. The terminal goal of the team is to converge on the source target.

4. A higher status value implies a higher probability that the source is located near the corresponding coordinates.

Alpha-beta agents are *eusocial* [Mcfarland 1994] in nature; agents must cooperate to succeed. Agents always broadcast their most current sensor data as a normative behavior. An agent's model of the environment is based solely on their current local sensory data and the current shared data obtained from the other members of the team. Individual agents have no sensor memory and consequently cannot locate a source alone. As such, the alpha-beta strategy is a reactive collective search strategy rather than a collaborative strategy. Agents are implicitly cooperative, and do not use explicit forms of collaboration. The alpha-beta strategy is a behavior-based control strategy closely related to the approach of Mataric [1994]. Alpha-beta teams behave in a manner similar to that of of simple insect societies [Kube and Zhang 1993]. Alpha-beta agents search without centralized leadership or hierarchical coordination. The primary collective mode of an alpha-beta team is to aggregate in a region of high-intensity, without any other objectives. Alpha-beta teams are robust to single-point fail-stop failures in team members; agents simply use the latest data transmitted by other team members without regard to the identity of the sender. Alpha-beta coordination requires a minimum of knowledge about the search environment. Agents have no prior assumptions about the nature of the intensity surface, its spatial coherence, gradient field, or any other analytical information. As such, the alpha-beta strategy is intended to be as general-purpose and as assumption-free as possible. In formulating the alpha-beta strategy, we have carefully constructed the problem context and agent capabilities to focus the research in a particular direction, namely away from traditional symbolic AI approaches and towards the dynamical systems and behavior-based/emergent behavior approach. This paper presents the alpha-beta concept and exhibits preliminary simulation data in an ideal environment. Our goal is to demonstrate collective coordination based on self-selected dynamical control laws that change in response to the collective state of the team.

3 Alpha Beta Coordination Algorithms

A full mathematical treatment of alpha-beta coordination is in progress [Goldsmith & Robinett 1998b] but is beyond the scope of this paper. The current state-space formulation comprises a system of non-linear, time-varying difference equations of order N, where N is the instantaneous number of agents. The issues of primary importance are stability, energy efficiency, convergence, and steady-state localization error.

A simple social metaphor provides an intuitively satisfying if imprecise description of the basis for alpha-beta coordination algorithms. The cohesion of an alpha-beta society is based on a common normative goal: each agent is motivated to improve its social status by associating with other agents of higher status. Social status is determined by a scalar function of the shared sensor data communicated by other agents. The only assumption underlying alpha-beta algorithms is that the status function orders points in the search space according to the probability that a source is located at the point. On each decision cycle, each agent broadcasts it current social status as a scalar value, s_i, along with a location vector, $\mathbf{v_i}$, to all other agents, and receives their status values in return. An agent attempts to improve its standing through emulation by moving to a region occupied by agents reporting superior status. This simple goal pressures agents to: (1) aggregate into groups; and (2) to aggregate in the region of highest known status. To determine its next destination, each agent first computes the common ordered set $\mathbf{V}=\{\mathbf{v}_i\}$ according to the linear ordering ($\leq$) of agents provided by the status readings $\mathbf{S}=\{s_i\}$[4]. The agent uses **S** to partition its fellow agents into two castes. The *alpha caste* is the set $\mathbf{A_0}$ of all agent positions corresponding to agents that have a social standing superior to agent a_0: $\mathbf{A_0} = \{\mathbf{v}_k|s_k > s_0\}$. The *beta caste* $\mathbf{B_0}$ is the set of all agent positions corresponding to agents with lower social standing than agent a_0: $\mathbf{B_0} = \{\mathbf{v}_k|s_k > s_0\}$. The beta set $\mathbf{B_0}$ includes agents of equal status because an agent always seeks to improve its current status. There are a variety of approaches to using the alpha and beta sets to generate the agent's next heading. The vectors in the set $\mathbf{A_0}$ can be used to influence the agent to move towards its members, creating a social pressure to improve called *alpha-pull*. The vectors in the set $\mathbf{B_0}$ can be used to influence the agent to move away from its members, creating a second social pressure to improve called *beta-push*. Either set or **V** itself can be used in a variety of ways to provide pressure to aggregate. Alpha-pull and beta push are heuristic in nature and do not necessarily lead to average improvement in arbitrary environments. Designing and testing different decision rules based on the data vectors in **V**, $\mathbf{A_0}$, and $\mathbf{B_0}$, or subsets thereof, is the means for investigating the different global behaviors of alpha-beta teams.

A special case of importance is when $\mathbf{V}=\mathbf{A_i}=\mathbf{B_i}$. In this case every agent has identical status, corresponding to the zero-information (maximum information entropy) state previously mentioned. When a zero-information state is detected, the team can disperse to broaden the search area by using beta-push (all members are in the beta sets of all other members) to compute a trajectory that leads the agents on the outer edges of the cohort region away from the team's centroid. As the density of the team decreases, more agents are free to move away from the centroid, eventually resulting in a dispersed team. A minimum limit on team density prevents the ultimate loss of team coherence. If the team members cannot find the sensate region, they must resort

[4] The unordered set of of readings can be used to compute the obvious non-uniform gradient estimates. We have investigated gradient search algorithms and use them as a baseline for comparison of alpha-beta performance. Some forms of alpha-beta algorithms currently under investigation use gradient estimates for alpha decisions.

to a collaborative search mode as mentioned previously.

If **V**=Ø, the agent is alone. For the purposes of this research, agents that lose contact with the team remain immobilized. This “hug a tree” philosophy saves energy but may not lead to a reunion with the team and to eventual arrival at the target source. A variety of possible solo behaviors will be investigated later, including random search, gradient search, and using the last known heading to determine the agent’s trajectory.

The general form of the alpha-beta update rule uses a linear combination of the vector data in **V**:

$$\mathbf{v}_i(k+1) = \mathbf{v}_i(k) + \mathbf{a}(k)[\mathbf{v}(k) - \mathbf{v}_i(k)] \qquad (1)$$

where **a** is a weighting vector derived from the application of some scalar function to the current status measurements **S** corresponding to the vectors in **v**. The nature of the function applied to **S** and the specific subset of vectors selected from **V** determine the group behavior exhibited by this version of alpha-beta teaming.

The alpha set **A** contains a distinguished subset of elements: the agent or agents with the highest status value. An agent with the highest status in the cohort has no alpha caste; **A**=ø. These agents are the *ø-alpha* agents and cannot experience alpha-pull. The choice of a decision rule for a *ø-alpha* agent is limited two possibilities: (1) don’t move; and (2) move away from the team along a vector derived from the B-vectors (beta-push). In the first option, the ø-alpha[5] identifies the location of highest known status and acts as a stationary beacon for the rest of the team. This is a conservative strategy that saves energy and ensures that the agent remains at the top of the heap, but does not immediately explore the region around the highest intensity reading. The second option uses some form of beta-push to move the ø-alpha away from the team. This is a risky strategy because the status of the ø-alpha may decrease, but it provides more information to the team and can possibly shorten convergence time.

The beta set $\mathbf{B_0}$ also contains a distinguished subset of elements: the agent or agents with the lowest status value. These *ø-beta* agents represent the social floor of the team, and always use some form of alpha-pull to improve their status.

The remaining members of the cohort have non-empty alpha and beta sets. Such an agent can experience the effects of both alpha-pull and beta-push. There are many possible decision rules for determining the next heading based on the partition $\{\mathbf{A_0}, \mathbf{B_0}\}$. In general, an agent must decide whether to be radical or conservative in its attempt to improve its status. The approach taken here is to provide three classes of behavior. For an agent team with N agents the update rules are:

1. The ø-alpha agents use the conservative decision mode and remain immobile: $\mathbf{v}_i(k+1) = \mathbf{v}_i(k)$.

2. The m agents in **V** with the highest status values self-select alpha behavior and use

[5] Although there may be more than one ø-alpha , we use the singular hereafter.

the following update rule: $\mathbf{v}_i(k+1) = \mathbf{v}_i(k) + u[\mathbf{v}^*(k) - \mathbf{v}_i(k)]$, where $\mathbf{v}^*(k)$ is the location of a ø-alpha agent, selected at random, and u is a factor that provides pressure to move beyond the alpha agent along a line passing through the points $\mathbf{v}^*(k)$ and $\mathbf{v}_i(k)$. Note that $u > 1.0$ must hold for improvement.

3. The remaining N-m agents in **V** self-select beta behavior and use the following update rule: $\mathbf{v}_i(k+1) = \mathbf{v}_i(k) + \mathbf{a}(k)[\mathbf{v}(k) - \mathbf{v}_i(k)]$, where $\mathbf{v}(k)$ are all members of $\mathbf{A_i}$, and $\mathbf{a}(k)$ is the corresponding vector with elements $a_j = s_j/D$, and

$$D = \sum s_k,\ k=1,\ N \qquad (2)$$

Under this regime, self-selected alpha agents attempt to exceed the performance of the stationary ø-alpha agent by attempting to overrun it. Self-selected beta agents compute a weighted average of the alpha vectors based on normalized status values and move towards the resultant. A conservative beta agent seeks to improve its status to the average status of its corresponding alpha set by moving to the point of the center-of-mass of the alpha set. This "safety in numbers" approach provides a tendency to aggregate in the most current region of highest known performance, but averages many alpha status positions to reduce noise and the influence of outliers. This behavior provides the beta population with some inertia, but still retains the tendency to improve the status of the population on average.

The important parameters in this regime are u, the "overrun factor" that determines the amount by which an alpha will attempt to move beyond a ø-alpha agent, and the alpha ratio, defined as =m/N, that determines the proportion of alpha agents exploring the search space.

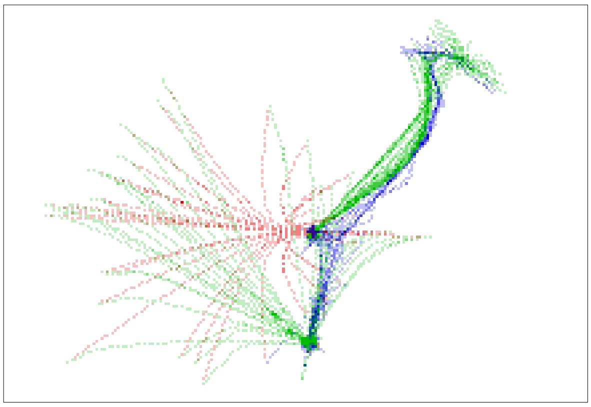

Figure 1: A team of 50 agents start in the upper right and locate a source at the center of the figure. The source intensity drops to zero and agents disperse to the right to locate another source. The source reappears in the lower center and agents once again converge upon it.

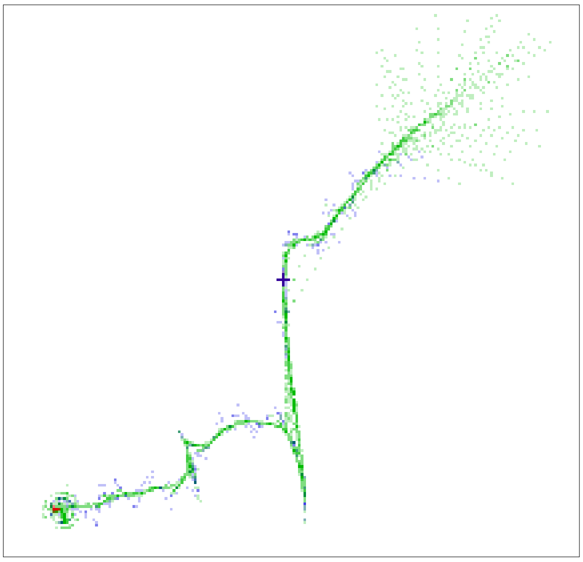

Figure 2: A team of 50 agents start in the upper right and eventually locate a source at the lower left. The annular region around the source results from alpha agents continuously searching around the source. The search trajectory is typical of an alpha-beta agent team.

4 Simulations and Results

The alpha-beta coordination strategy was simulated in an ideal 2-D world using ideal agents. The intent of these initial simulations was to study the convergence and coordination properties of alpha-beta rather than evaluate alpha-beta in a realistic environment. The simulations provide a best-case baseline against which various complications such as communications noise and sensor noise can be evaluated later on. The world is free of obstacles, ambient noise, communications errors and convection currents that make the source intensity field non-stationary and time-varying. Ideal agents are point-masses with no area, so crowding is not an issue. Ideal agents have noise-free sensors, and movement on each step is bounded.

The target source was a radial emitter with exponential decay factor b and a uniform random noise component w:

$$Z(r) = w + \exp -(r \cdot b) \qquad (3)$$

where r is the radial distance from the origin. The metric of interest for this study is the mean-squared distance from the target, a measure of the team's learning rate and steady-state convergence error. For each simulation run, alpha-beta agents are initially positioned with the same distribution in the x-y plane. A control group comprising agents with identical starting points but with knowledge of the source location provide a baseline for

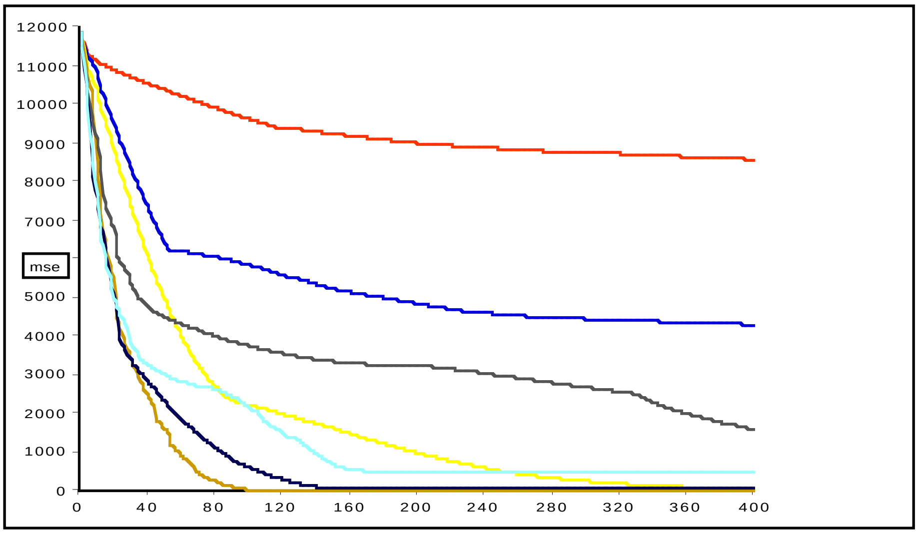

Figure 3. Mean-squared error vs. step for u=2.0 and (1) R=0.1 (upper); (2) R=0.2; (3) R=0.4; (4)R=0.5(lower); (6) R=0.6;(7)R=0.8; (8)R=1.0(third from top). Convergence is for ß=0.5. Notice the diminishing returns for ß>0.5.

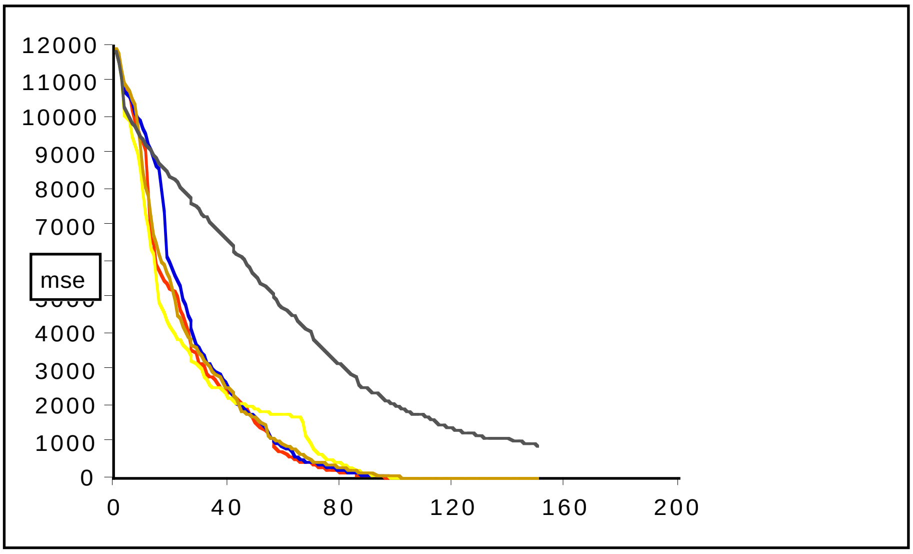

Figure 4: Mean-squared error vs. step for R=0.5 and: (1) u=2.0; (2) u=3.0 ; (3) u=4.0;(4) u=8.0;(5) u=10.0(upper).

learning curve for the team. Figures 1 and 2 show typical traces of alpha-beta agents.

The simulation results confirm that the team can find a source under ideal conditions.

The alpha ratio is R critical to effective search. A critical mass of alpha agents is needed to influence the beta agents to follow the alpha trajectory. A ratio of not less than 0.4 is needed for reliable search given a u value of 2.0. Maximum convergence rate and minimum steady-state mean-squared error occur at R=0.5. Figure 3 shows the learning curves for various alpha-beta ratios. Convergence rate is somewhat sensitive to the alpha u parameter as expected, favoring greater values of u at the expense of increased steady-state mean squared error. Very large values of u slow the convergence rate and lead to larger steady-state errors.

5 Related Work

The emergence of global behavior from local interactions among autonomous agents has been studied extensively. Investigations of collective behavior in robots are considerably more rarefied, and studies involving collective search are rarer still. The foraging problem [Arkin & Hobbs 1993, Goss & Deneubourg 1992, Mataric 1994, Steels 1990], in which robots collect objects scattered in the environment, is a canonical problem related to the source location problem.

The alpha-beta strategy falls squarely in the behavior-based control camp [Brooks 1991, Brooks 1986, Mataric 1992]. Mataric (1994;1995) describes group behaviors in terms of combinations of basis behaviors invoked by sensor inputs. Flocking, a commonplace group behavior, comprises the primitive basis behaviors of safe-wandering, homing, aggregation, and dispersion. Following and aggregation make up surrounding, and herding is composed of surrounding and flocking. Flocking, homing, following, aggregation, and dispersion are all behaviors that arise under alpha-beta coordination, but are not accomplished by compositions of explicitly programmed basis behaviors. Different behaviors are obtained in alpha-beta coordination through variations on the update equation (1).Goldberg and Mataric [1997] describe pack and caste criteria for partitioning a robot team to achieve arbitration of spatial interference. Their approach shares with alpha-beta the concept of behavioral switching based on the collective state. The dynamics-based approach [Large, Henrik, and Bajcsy 1997] is also similar to alpha-beta in its use of of a vector-based dynamical system to generate robot behaviors. Social entropy, a measure of the behavioral diversity in a robot team based on information entropy, has been presented in [Balch 1997]. This is a potential metric for alpha-beta regimes and we will investigate its application in future research.

6 Discussion and Future Work

We have demonstrated the concept of dynamic social partitioning as a means to provide collective benefits to an agent team searching for source targets. Initial simulations confirm the ability of the team to find a source and stabilize the steady-state mean-squared error.

Our future research will focus on further investigations of alternative forms of alpha-beta algorithms inspired by molecular dynamics and statistical mechanics. We intend to investigate new forms of interaction rules that are based on non-linear functions of the entire measurement set rather than on partitions of the measurements. We will also investigate dynamic adjustment of the alpha u parameter and the alpha/beta ratio through reinforcement learning techniques under the alpha-beta regime presented in this paper. Simulations involving more realistic environments containing obstacles, convection effects on chemical plumes, and more detailed models of robotic vehicles will be conducted on parallel processors for large numbers of agents if required. Ultimately, we will attempt to implement alpha-beta strategy on actual robotic vehicles.

Acknowledgements

We wish to thank our colleagues at the Advanced Information Systems Laboratory, Laurence Phillips and Shannon Spires, for technical discussions and comments on this paper, and for coding and debugging portions of the simulation.

References[6]

Arkin, R. and Hobbs, J. 1994. Dimensions of communication and social organization in multi-agent robotic systems. In *Proc. Simulation of Adaptive Behavior.*

Balch, T. 1997. Social entropy: A new metric for learning multi-robot teams. In *Proc. FLAIRS-97*.

Brooks, R. 1986. A robust layered control system for a mobile robot. *IEEE Journal of Robotics and Automation* RA-2.

Brooks, R. 1991. Intelligence without reason. In *Proc. IJCAI-91*.

Cao, U., Fukunaga, A., Kahng, A., and Meng, F. 1995. Cooperative mobile robotics: Antecedents and directions. In *Proc. of IEEE/RSJ IROS.*

Goldberg, D. and Mataric, M. 1997. Interference as a tool for designing and evaluating multi-robot controllers. In *Proceedings, AAAI-97, Providence, Rhode Island, July 27-31, 1997, 637-642*.

Goldsmith, S. and Robinett, R. 1998a. Collaborative search by mobile robots- Part I: Problem definition. In progress.Preprint at www.sandia.gov/aisl/robotics/collaboration

Goldsmith, S., and Robinett, R. 1998b. Stability and performance of alpha-beta coordination. Working paper. Preprint: www.sandia.gov/aisl/robotics/ab-stability

Goss, S., and Deneubourg, J. 1992. Harvesting by a group of robots. In *Proc. European Conf. on Artificial Life*.

Kube, C. and Zhang, H. 1993. Collective robotics: From social insects to robots.

[6] The AISL website will be opened to the Internet before the July 4 conference opening.

Adaptive Behavior, 2(2).

Klarer, P. 1998. Flocking small smart machines: An experiment in cooperative multi-machine control. Sandia National Laboratories Report. In progress.

Lamport, L., Shostak, R., and Pease, M. ``The Byzantine Generals Problem," A*CM Transactions on Programming Languages and Systems 4, 3* (July 1982), 382--401.

Large, E., Henrik, C., Bajcsy, R. 1997. Dynamic Robot Planning: cooperation through competition. In *Proceedings of the IEEE International Conference on Robotics and Automation,* Albuquerque, NM.

Mataric, M. 1992. Behavior-based systems: Key properties and implications. In *IEEE International Conference on Robotics and Automation, Workshop on Architectures for Intelligent Control Systems*.

Mataric, M. 1994. Interaction and intelligent behavior. *MIT AI Lab Tech report AITR-1495*.

Mataric, M. 1995. Designing and understanding adaptive group behavior. *Adaptive Behavior (4:1)* .

McFarland, D. 1994. Towards robot cooperation. In *Proc. Simulation of Adaptive Behavior*.

Spires, S. , and Goldsmith, S. 1998. Exhaustive Geographic Search with Mobile Robots Along Space-Filling Curves. Submitted to Collective Robotics Workshop, Agent World '98.

Steels, L. 1990. Cooperation between distributed agents through self-organization. In *European Workshop on Modeling Autonomous Agents in a Multi-Agent World.*